

ABSTRACT OF THE DISCLOSURE

[1068] Typically, missing read operations instances account for a small fraction of the operations instances of an application, but for nearly all of the performance degradation due to access latency. Hence, a small predictor structure maintains sufficient information for performing value prediction for the small fraction of operations (the missing instances of read operations) that account for nearly all of the access latency performance degradation. With such a small predictor structure, a processor value predicts for selective instances of read operations, those selective instances being read operations that are unavailable in a first memory (e.g., those instances of read operations that miss in L2 cache). Respective actual values for prior missing instances of the read operations are stored and used for value predictions of respective subsequent instances of the read operations. The value predictions are, at least partially, based on accuracy of value predictions for prior corresponding missing instances of the read operations.